Evolving Agentic AI Architecture

1. The Core Stack: Building Robust Agent Systems

Participants detailed a sophisticated "full stack" for AI agents that moves beyond basic chat

Fundamental Components

A standard system requires a Large Language Model (LLM) for text generation, an iterative self-correction loop for "thinking," and access to external tools for taking actions.

Memory & Context Management

Current technology cannot reliably maintain long-term conversation tracks; therefore, every agent call should include fresh instructions and context managed by an external software orchestrator.

Deterministic Data Retrieval

To ensure safety in production, experts recommend using deterministic code (SQL/NoSQL) via function descriptions rather than letting agents generate their own queries.

Graph RAG & Ontology

Implementing Graph RAG (Retrieval-Augmented Generation) helps agents understand specific domains and restricts their actions to explicit permissions.

This summary captures the core technical and strategic insights from the "Agentic AI Architecture Evolving" discussion, held in February 2026 in connection with the AI Finland Meetup and facilitated by Futurice. The session focused on the transition from simple AI assistants to autonomous agents and the rigorous architectural controls required to deploy them safely.

Evolving Agentic AI Architecture

AI Finland Meetup Highlights Feb 2026

2. Control Mechanisms & Governance

As agents gain autonomy, the discussion emphasized the critical need for "guardrails" to prevent errors and endless loops:

Self-Steering Rule Sets

Agents should be designed to output their own rule criteria before execution; if a supervising system detects a rule breach, it "kills" the process to allow for self-correction.

The Gatekeeper Model

For memory, a "gatekeeper" process should filter and decide which information is truly relevant for future use to keep the system's memory tight and valuable.

Markdown-Driven Alignment

One highlighted technique uses simple markdown files to define brand voice, colors, and tone, allowing leadership to guide marketing execution through code-like definitions.

Human-in-the-Loop

Future workflows will likely see AI prioritizing the most critical changes for human review from massive queues, as humans cannot monitor every autonomous action.

3. Market Realities & The Road Ahead

The group addressed the current state of implementation, particularly within the Finnish and Nordic ecosystems:

Rapid Prototyping

AI agents are drastically compressing development time, allowing "proof of concept" tasks that once took weeks to be completed in a single afternoon.

The "Ambition Gap"

Observations from the meeting noted that Swedish investors' higher ambition and a larger funding ecosystem are currently accelerating development faster than in Finland.

Production Maturity

While some large companies have fully automated IT support via agents, most current deployments remain "assistant-style" features rather than end-to-end business functions.

Compliance & The AI Act

High-risk applications face significant hurdles in tracking agent decision-making chains for regulatory compliance.